

Energy and Sampling Constrained Asynchronous Communication

Aslan Tchamkerten, Venkat Chandar, and Giuseppe Caire

Abstract—The minimum energy, and, more generally, the minimum cost, to transmit one bit of information has been recently derived for bursty communication when information is available infrequently at random times at the transmitter. This result assumes that the receiver is always in the listening mode and samples all channel outputs until it makes a decision. If the receiver is constrained to sample only a fraction $\rho \in (0, 1]$ of the channel outputs, what is the cost penalty due to sparse output sampling?

Remarkably, there is no penalty: regardless of $\rho > 0$ the asynchronous capacity per unit cost is the same as under full sampling, *i.e.*, when $\rho = 1$. Moreover, there is not even a penalty in terms of decoding delay—the elapsed time between when information is available until when it is decoded. This latter result relies on the possibility to sample adaptively; the next sample can be chosen as a function of past samples. Under non-adaptive sampling, it is possible to achieve the full sampling asynchronous capacity per unit cost, but the decoding delay gets multiplied by $1/\rho$. Therefore adaptive sampling strategies are of particular interest in the very sparse sampling regime.

Index Terms—Asynchronous communication; bursty communication; capacity per unit cost; energy; error exponents; hypothesis testing; sequential decoding; sensor networks; sparse communication; sparse sampling; synchronization

I. INTRODUCTION

IN many emerging technologies, communication is sparse and asynchronous, but it is essential that when data is available, it is delivered to the destination as timely and reliably as possible. Examples are sensor networks monitoring rare but critical events, such as earthquakes, forest fires, or epileptic seizures.

For such settings, [1] characterized the asynchronous capacity per unit cost based on the following model.

This work was supported in part by an Excellence Chair Grant from the French National Research Agency (ACE project).

A. Tchamkerten is with the Department of Communications and Electronics, Telecom ParisTech, 75634 Paris Cedex 13, France. Email: aslan.tchamkerten@telecom-paristech.fr.

V. Chandar is with MIT Lincoln Laboratory, Lexington, MA 02420, USA. Email: vchandar@mit.edu.

G. Caire is with the Viterbi School of Engineering, University of Southern California, Los Angeles, USA. Email: caire@usc.edu.

There are B bits of information that are made available to the transmitter at some random time ν , and need to be communicated to the receiver. The B bits are coded and transmitted over a memoryless channel using a sequence of symbols that have costs associated with them. The rate R per unit cost is the total number of bits divided by the cost of the transmitted sequence. Asynchronism is captured here by the fact that the random time ν is not known *a priori* to the receiver. However both transmitter and receiver know that ν is distributed (*e.g.*, uniformly) over a time horizon $[1, \dots, A]$. At all times before and after the actual transmission, the receiver observes “pure noise.” The noise distribution corresponds to a special input “idle symbol” \star being sent across the channel (for example, in the case of a Gaussian channel, this would be the 0, *i.e.*, no transmit signal).

The goal of the receiver is to reliably decode the information bits by sequentially observing the outputs of the channel.

A main result in [1] is a single-letter characterization of the asynchronous capacity per unit cost $C(\beta)$ where

$$\beta \stackrel{\text{def}}{=} \frac{\log A}{B}$$

denotes the *timing uncertainty per information bit*. While this result holds for arbitrary discrete memoryless channels and arbitrary input costs, the underlying model assumes that the receiver is always in the listening mode: every channel output is observed until decoding happens.

What happens when the receiver is constrained to observe a fraction $0 < \rho \leq 1$ of the channel outputs? In this paper, it is shown that the asynchronous capacity per unit cost is *not* impacted by a sparse output sampling. More specifically, the asynchronous capacity per unit cost satisfies

$$C(\beta, \rho) = C(\beta, 1)$$

for any asynchronism level $\beta > 0$ and sampling frequency $0 < \rho \leq 1$. Moreover, the decoding delay is minimal: the elapsed time between when information starts being sent and when it is decoded is the same as under full sampling. This result uses the possibility for the receiver to sample adaptively: the next sample can be chosen as a function of past observed samples. In

fact, under non-adaptive sampling, it is still possible to achieve the full sampling asynchronous capacity per unit cost, but the decoding delay gets multiplied by a factor $1/\rho$ or $(1 + \rho)/\rho$ depending on whether or not \star can be used for code design. Therefore, adaptive sampling strategies are of particular interest in the very sparse regime.

We end this section with a brief review of studies related to the above communication model. This model was introduced in [2], [3]. Both of these works focused mainly on the *synchronization threshold*—the largest level of asynchronism under which it is still possible to communicate reliably. In [3], [4] communication rate is defined with respect to the decoding delay, the expected elapsed time between when information is available and when it is decoded. Capacity upper and lower bounds are established and shown to be tight for certain channels. In [4] it is also shown that so-called training-based schemes, where synchronization and information transmission use separate degrees of freedom, need not be optimal in particular in the high rate regime.

The finite message regime has been investigated by Polyanskiy in [5] when capacity is defined with respect to the codeword length, *i.e.*, same setting as [1] but with unit cost per transmitted symbol. A main result in [5] is that dispersion—a fundamental quantity that relates rate and error probability in the finite block length regime—is unaffected by the lack of synchronization. Whether or not this remains true under sparse output sampling is an interesting open issue.

Note that the seemingly similar notions of rates investigated in [3], [4] and [1], [5] are in fact very different. In particular, capacity with respect to the expected decoding delay remains in general an open problem.

A “slotted” version of the above communication model was considered in [6] by Wang, Chandar, and Wornell where communication now can happen only in one of consecutive slots of the size of a codeword. For this model, the authors investigated the tradeoff between the false-alarm event (the decoder declares a message before even it is sent) and the miss event (the decoder misses the sent codeword).

The previous works consider point-to-point communication. A (diamond) network configuration was recently investigated by Shomorony, Etkin, Parvaresh, and Alvestimehr in [7] who provided bounds on the minimum energy needed to convey one bit of information across the network.

In above models, although communication is bursty, information transmission is contiguous since it always lasts the codeword duration. A complementary setup proposed by Khoshnevisan and Laneman [8] considers a

bursty communication scenario caused by an intermittent codeword transmission. This model can be seen as a slotted variation of the purely insertion channel model, the latter being a particular case of the general insertion, deletion, and substitution channel introduced by Dobrushin [9].

This paper is organized as follows. Section II contains some background material and extends the model developed in [1] to allow for sparse output sampling. Section III contains the main results and briefly discusses extensions to a decoder-universal setting and to a multiple access setup. Finally Section IV is devoted to the proofs.

II. MODEL AND PERFORMANCE CRITERION

The asynchronous communication model we consider captures the following general features:

- Information is available at the transmitter at a random time;
- The transmitter can choose when to start sending information based on when information is available and based on what message needs to be transmitted;
- There is a cost associated to each channel input;
- Outside the information transmission period the transmitter stays idle and the receiver observes noise;
- The decoder is sampling constrained and can observe only a fraction of the channel outputs.
- Without knowing a priori when information is available, the decoder should decode reliably and as early as possible, on a sequential basis.

The model is now specified. Communication is discrete-time and carried over a discrete memoryless channel characterized by its finite input and output alphabets

$$\mathcal{X} \cup \{\star\} \quad \text{and} \quad \mathcal{Y},$$

respectively, and transition probability matrix

$$Q(y|x),$$

for all $y \in \mathcal{Y}$ and $x \in \mathcal{X} \cup \{\star\}$. The alphabet \mathcal{X} may or may not include \star . Without loss of generality, we assume that for all $y \in \mathcal{Y}$ there is some $x \in \mathcal{X} \cup \{\star\}$ for which $Q(y|x) > 0$.

Given $B \geq 1$ information bits to be transmitted, a codebook \mathcal{C} consists of

$$M = 2^B$$

codewords of length $n \geq 1$ composed of symbols from \mathcal{X} .

A randomly and uniformly chosen message m arrives at the transmitter at a random time ν , independent of

m , and uniformly distributed over $[1, \dots, A]$, where the integer

$$A = 2^{\beta B}$$

characterizes the *asynchronism level* between the transmitter and the receiver, and where the constant

$$\beta \geq 0$$

denotes the *timing uncertainty per information bit*, see Fig. 1.

We consider one-shot communication, *i.e.*, only one message arrives over the period $[1, 2, \dots, A]$. If $A = 1$, the channel is said to be synchronous.

Given ν and m , the transmitter chooses a time $\sigma(\nu, m)$ to start sending codeword $c^n(m) \in \mathcal{C}$ assigned to message m . Transmission cannot start before the message arrives or after the end of the uncertainty window, hence $\sigma(\nu, m)$ must satisfy

$$\nu \leq \sigma(\nu, m) \leq A \quad \text{almost surely.}$$

In the rest of the paper, we suppress the arguments ν and m of σ when these arguments are clear from context.

Before and after the codeword transmission, *i.e.*, before time σ and after time $\sigma + n - 1$, the receiver observes “pure noise,” Specifically, conditioned on the event $\{\nu = t\}$, $t \in \{1, \dots, A\}$, and on the message to be conveyed m , the receiver observes independent channel outputs

$$Y_1, Y_2, \dots, Y_{A+n-1}$$

distributed as follows. For

$$1 \leq i \leq \sigma(t, m) - 1$$

or

$$\sigma(t, m) + n \leq i \leq A + n - 1,$$

the Y_i 's are “pure noise” symbols, *i.e.*,

$$Y_i \sim Q(\cdot | \star).$$

For $\sigma \leq i \leq \sigma + n - 1$

$$Y_i \sim Q(\cdot | c_{i-\sigma+1}(m))$$

where $c_i(m)$ denotes the i th symbol of the codeword $c^n(m)$.

The receiver operates according to a sampling strategy and a sequential decoder. A sampling strategy consists of “sampling times” which are defined as an ordered collection of random time indices

$$\mathcal{S} = \{(S_1, \dots, S_\ell) \subseteq \{1, \dots, A+n-1\} : S_i < S_j, i < j\}$$

where S_j is interpreted as the j th sampling time.

The sampling strategy is either non-adaptive or adaptive. It is non-adaptive when the sampling times given




Fig. 1. Time representation of what is sent (upper arrow) and what is received (lower arrow). The “ \star ” represents the “idle” symbol. Message m arrives at time ν and starts being sent at time σ . The receiver samples at the (random) times S_1, S_2, \dots and decodes at time S_τ based on τ output samples.

by \mathcal{S} are all known before communication starts, hence \mathcal{S} is independent of Y_1^{A+n-1} . The strategy is adaptive when the sampling times are function of past observations. This means that S_1 is an arbitrary value in $\{1, \dots, A + n - 1\}$, possibly random but independent of Y_1^{A+n-1} and, for $j \geq 2$,

$$S_j = g_j(\{Y_{S_i}\}_{i < j})$$

for some (possibly randomized) function

$$g_j : \mathcal{Y}^{j-1} \rightarrow \{S_{j-1} + 1, \dots, A + n - 1\}.$$

Notice that ℓ , the total number of output samples, may be random under adaptive sampling, but also under non-adaptive sampling, since the strategy may be randomized (but still independent of the channel outputs Y_1^{A+n-1}).

Once the sampling strategy is fixed, the receiver decodes by means of a sequential test (τ, ϕ) , where τ , the decision time, is a stopping time with respect to the sampled sequence

$$Y_{S_1}, Y_{S_2}, \dots$$

indicating when decoding happens,¹ and where ϕ is the decoding function, *i.e.*, a map

$$\phi : \mathcal{O} \rightarrow \{1, 2, \dots, M\}$$

where

$$\mathcal{O} \stackrel{\text{def}}{=} \{Y_{S_1}, Y_{S_2}, \dots, Y_{S_\tau}\}$$

is the set of observed samples. Hence, decoding happens at time S_τ on the basis of τ output samples. Since there are at most $A + n - 1$ sampling times, τ is bounded by $A + n - 1$.

A code $(\mathcal{C}, \mathcal{S}, (\tau, \phi))$ is defined as a codebook, a receiver sampling strategy, and a decoder (decision time

¹Recall that a (deterministic or randomized) stopping time τ with respect to a sequence of random variables Y_1, Y_2, \dots is a positive, integer-valued, random variable such that the event $\{\tau = t\}$, conditioned on the realization of Y_1, Y_2, \dots, Y_t , is independent of the realization of Y_{t+1}, Y_{t+2}, \dots , for all $t \geq 1$.

and decoding function). Throughout the paper, whenever clear from context, we often refer to a code using the codebook symbol \mathcal{C} only, leaving out an explicit reference to the sampling strategy and to the decoder.

Definition 1 (Error probability). The maximum (over messages) decoding error probability of a code \mathcal{C} is defined as

$$\mathbb{P}(\mathcal{E}|\mathcal{C}) \stackrel{\text{def}}{=} \max_m \frac{1}{A} \sum_{t=1}^A \mathbb{P}_{m,t}(\mathcal{E}_m), \quad (1)$$

where the subscripts “ m, t ” denote conditioning on the event that message m arrives at time $\nu = t$, and where \mathcal{E}_m denotes the error event that the decoded message does not correspond to m , *i.e.*,

$$\mathcal{E}_m \stackrel{\text{def}}{=} \{\phi(\mathcal{O}) \neq m\}.$$

Definition 2 (Cost of a Code). The (maximum) cost of a code \mathcal{C} with respect to a cost function $k : \mathcal{X} \rightarrow [0, \infty]$ is defined as

$$K(\mathcal{C}) \stackrel{\text{def}}{=} \max_m \sum_{i=1}^n k(c_i(m)).$$

Assumption: throughout the paper we make the assumption that the only possible zero cost symbol is \star . When $\star \in \mathcal{X}$ the transmitter can stay idle at no cost. When $\star \notin \mathcal{X}$ then $k(x) > 0$ for any $x \in \mathcal{X}$, which captures the situation where a “standby” mode may not be possible at zero cost. The other cases—investigated in [1] under full sampling—are either trivial (when \mathcal{X} contains two or more zero costs symbols) or arguably unnatural (\mathcal{X} contains a zero cost symbol that differs from \star or when $\star \in \mathcal{X}$ and all \mathcal{X} contains only nonzero cost symbols).

Below, \mathbb{P}_m denotes the output distribution conditioned on the sending of message m . Hence, by definition we have

$$\mathbb{P}_m(\cdot) \stackrel{\text{def}}{=} \frac{1}{A} \sum_{t=1}^A \mathbb{P}_{m,t}(\cdot).$$

Definition 3 (Sampling Frequency of a Code). Given $\varepsilon > 0$, the sampling frequency of a code \mathcal{C} , denoted by $\rho(\mathcal{C}, \varepsilon)$, is the relative number of channel outputs that are observed until a message is declared. Specifically, it is defined as the smallest $r \geq 0$ such that

$$\min_m \mathbb{P}_m(\tau/S_\tau \leq r) \geq 1 - \varepsilon.$$

(Recall that S_τ refers to the last sampling time.)

Definition 4 (Delay of a Code). Given $\varepsilon > 0$, the (maximum) delay of a code \mathcal{C} , denoted by $d(\mathcal{C}, \varepsilon)$, is defined as the smallest integer l such that

$$\min_m \mathbb{P}_m(S_\tau - \nu \leq l - 1) \geq 1 - \varepsilon.$$

We now define capacity per unit cost under the constraint that the receiver has access only to a limited number of channel outputs:

Definition 5 (Asynchronous Capacity per Unit Cost under Sampling Constraint). \mathbf{R} is an achievable rate per unit cost at timing uncertainty per information bit β and sampling frequency ρ , if there exists a sequence of codes $\{\mathcal{C}_B\}$ and a sequence of positive numbers ε_B with $\varepsilon_B \xrightarrow{B \rightarrow \infty} 0$ such that for all B large enough

- 1) \mathcal{C}_B operates at timing uncertainty per information bit β ;
- 2) the maximum error probability $\mathbb{P}(\mathcal{E}|\mathcal{C}_B)$ is at most ε_B ;
- 3) the rate per unit cost

$$\frac{B}{K(\mathcal{C}_B)}$$

is at least $\mathbf{R} - \varepsilon_B$;

- 4) the sampling frequency satisfies $\rho(\mathcal{C}_B, \varepsilon_B) \leq \rho + \varepsilon_B$;
- 5) the delay satisfies²

$$\frac{1}{B} \log(d(\mathcal{C}_B, \varepsilon_B)) \leq \varepsilon_B.$$

Notice that the last requirement asks for a subexponential delay.

The asynchronous capacity per unit cost, denoted by $\mathbf{C}(\beta, \rho)$, is the supremum of achievable rates per unit cost.

Two basic observations:

- $\mathbf{C}(\beta, \rho)$ is a non-increasing function of β for fixed ρ ;
- $\mathbf{C}(\beta, \rho)$ is a non-decreasing function of ρ for fixed β .

In particular, for any fixed $\beta \geq 0$

$$\max_{\rho \geq 0} \mathbf{C}(\beta, \rho) = \mathbf{C}(\beta, 1).$$

Capacity per unit cost under full sampling $\mathbf{C}(\beta, 1)$ is characterized in the following theorem:

Theorem 1 ([1] Theorem 1). *For any $\beta \geq 0$*

$$\mathbf{C}(\beta, 1) = \max_X \min \left\{ \frac{I(X; Y)}{\mathbb{E}[k(X)]}, \frac{I(X; Y) + D(Y||Y_\star)}{\mathbb{E}[k(X)](1 + \beta)} \right\}, \quad (2)$$

where \max_X denotes maximization with respect to the channel input distribution P_X , where $(X, Y) \sim P_X(\cdot)Q(\cdot|\cdot)$, where Y_\star denotes the random output of the channel when the idle symbol \star is transmitted (*i.e.*,

²Throughout the paper \log is always to the base 2.

$Y_\star \sim Q(\cdot|\star)$, where $I(X;Y)$ denotes the mutual information between X and Y , and where $D(Y||Y_\star)$ denotes the divergence (Kullback-Leibler distance) between the distributions of Y and Y_\star .³ ■

Let P_{X^\star} be a capacity per unit cost achieving input distribution, i.e., X^\star achieves the maximum in (2). As shown in the converse of the proof of [1, Theorem 1], codes that achieve the capacity per unit cost can be restricted to codes of (asymptotically) constant composition P_{X^\star} . Specifically, we have

$$\frac{B}{n_B(P_{X^\star})\mathbb{E}[k(X^\star)]} = C(\beta, 1)(1 - o(1)) \quad (B \rightarrow \infty)$$

where $n_B(P_{X^\star})$ denotes the length of the P_{X^\star} -constant composition codes achieving $C(\beta, 1)$. Now define

$$n_B^* \stackrel{\text{def}}{=} \min_{P_{X^\star}} n_B(P_{X^\star}) = \min_{X \in \mathcal{P}} \frac{B}{C(\beta, 1)\mathbb{E}[k(X)]}$$

where

$$\mathcal{P} \stackrel{\text{def}}{=} \{X : X \text{ achieves the maximum in (2)}\}.$$

From the achievability and converse of [1, Theorem 1], $\{n_B^*\}$ represent the smallest achievable delays for codes $\{\mathcal{C}_B\}$ achieving the asynchronous capacity per unit cost under full sampling $C(\beta, 1)$ in the sense that

$$d(\mathcal{C}_B, \varepsilon_B) \geq n_B^*(1 - o(1)) \quad (B \rightarrow \infty)$$

for any $\varepsilon_B \rightarrow 0$ as $B \rightarrow \infty$.

Our results, stated in the next section, say that the capacity per unit cost under sampling frequency $0 < \rho < 1$ is the same as under full sampling, i.e., $\rho = 1$. To achieve this, non-adaptive sampling is sufficient. However, if we also want to achieve minimum delay, then adaptive sampling is necessary. In fact, non-adaptive sampling strategies that achieve capacity per unit cost have a delay that grows at least as

$$\frac{n_B^*}{\rho}$$

or

$$\frac{n_B^*(1 + \rho)}{\rho}$$

depending on whether or not $\star \in \mathcal{X}$.

We end this section with a few notational conventions. We use $\mathcal{P}^\mathcal{X}$ to denote the set of distributions over the finite alphabets \mathcal{X} . Recall that the type of a string $x^n \in \mathcal{X}^n$, denoted by \hat{P}_{x^n} , is the probability over \mathcal{X} that assigns to each $a \in \mathcal{X}$ the number of occurrences of a within x^n divided by n [10, Chapter 1.2]. For instance, if $x^3 = (0, 1, 0)$, then $\hat{P}_{x^3}(0) = 2/3$ and $\hat{P}_{x^3}(1) = 1/3$. The joint

type \hat{P}_{x^n, y^n} induced by a pair of strings $x^n \in \mathcal{X}^n, y^n \in \mathcal{Y}^n$ is defined similarly. The set of strings of length n that have type P is denoted by \mathcal{T}_P^n . The set of all types over \mathcal{X} of strings of length n is denoted by $\mathcal{P}_n^\mathcal{X}$. Finally, we use $\text{poly}(\cdot)$ to denote a function that does not grow or decay faster than polynomially in its argument.

Throughout the paper we use the standard “big-O” Landau notation to characterize growth rates (see, e.g., [11, Chapter 3]).

III. RESULTS

In the sequel we denote by $C_a(\beta, \rho)$ and $C_{na}(\beta, \rho)$ the capacity per unit cost when restricted to adaptive and non-adaptive sampling, respectively.

Our first result characterizes the capacity per unit cost under non-adaptive sampling.

Theorem 2 (Non-adaptive sampling). *Under non-adaptive sampling it is possible to achieve the full-sampling capacity per unit cost, i.e.*

$$C_{na}(\beta, \rho) = C(\beta, 1) \quad \text{for any } \beta > 0, \rho > 0.$$

Furthermore codes $\{\mathcal{C}_B\}$ that achieve rate $\gamma C(\beta, 1)$, $0 \leq \gamma \leq 1$, satisfy

$$\lim_{\gamma \rightarrow 1} \liminf_{B \rightarrow \infty} \frac{d(\mathcal{C}_B, \varepsilon_B)}{n_B^*} \geq \frac{1}{\rho}$$

when $\star \in \mathcal{X}$, and satisfy

$$\lim_{\gamma \rightarrow 1} \liminf_{B \rightarrow \infty} \frac{d(\mathcal{C}_B, \varepsilon_B)}{n_B^*} \geq \frac{1 + \rho}{\rho}$$

when $\star \notin \mathcal{X}$. Finally, the above delay bounds are tight: for any $\varepsilon > 0$ and γ close enough to 1 there exists $\{\mathcal{C}_B\}$ and $\varepsilon_B \rightarrow 0$ as $B \rightarrow \infty$ such that

$$\liminf_{B \rightarrow \infty} \frac{d(\mathcal{C}_B, \varepsilon_B)}{n_B^*} \leq \frac{1}{\rho} + \varepsilon$$

for the case $\star \in \mathcal{X}$, and similarly for the case $\star \notin \mathcal{X}$.

Hence, even with a negligible fraction of the channel outputs it is possible to achieve the full-sampling capacity per unit. However, this comes at the expense of delay which gets multiplied by a factor $1/\rho$ or $(1 + \rho)/\rho$ depending on whether or not \star can be used for code design. This disadvantage is overcome by adaptive sampling:

Theorem 3 (Adaptive sampling). *Under adaptive sampling it is possible to achieve the full-sampling capacity per unit cost, i.e.*

$$C_{na}(\beta, \rho) = C(\beta, 1) \quad \text{for any } \beta > 0, \rho > 0.$$

³ Y_\star can be interpreted as “pure noise.”

Moreover, there exists $\{\mathcal{C}_B\}$ and $\varepsilon_B \rightarrow 0$ as $B \rightarrow \infty$ such that

$$d(\mathcal{C}_B, \varepsilon_B) = n_B^*(1 + o(1)).$$

The first part of Theorem 3 immediately follows from the first part of Theorem 2 since the set of adaptive sampling strategies include the set of non-adaptive sampling strategies. The interesting part of Theorem 3 is that adaptive sampling strategies guarantee minimal delay *regardless* of the sampling rate ρ , as long as it is non-zero.

What is an optimal adaptive sampling strategy? Intuitively, such a strategy should sample sparsely, with a sampling frequency of no more than ρ , under pure noise—for otherwise the sampling constraint is violated. It should also sample the entire sent codeword, and so densely sample during message transmission—for otherwise a rate per unit cost penalty is incurred. The main characteristic of a good adaptive sampling strategy is the criterion under which the sampling mode switches from sparse to dense. If the criterion is too conservative, *i.e.*, if the probability of switching under pure noise is too high, we might sample only part of the codeword, thereby incurring a cost loss. By contrast, if this probability is too low, we might not be able to accommodate the desired sampling frequency.

The proposed asymptotically optimal sampling/decoding strategy operates as follows—details are deferred to the proof of Theorem 3.

The strategy starts in the sparse mode, taking samples at times $S_j = \lceil j/\rho \rceil$, $j = 1, 2, \dots$. At each S_j , the receiver computes the empirical distribution (or type) of the last $\log(n)$ samples. If the probability of observing this type under pure noise is greater than $1/n^2$, the mode is kept unchanged and we repeat this test at the next round $j + 1$. Instead, if it is smaller than $1/n^2$, then we switch to the dense sampling mode, taking samples continuously for at most n time steps. At each of these steps the receiver applies a standard typicality decoding based on the past n output samples. If no codeword is typical with the channel outputs after these n times steps, sampling is switched back to the sparse mode. As it turns out, the threshold $1/n^2$ can be replaced by any decreasing function of n that decreases at least as fast as $1/n^2$ but not faster than polynomially in n .

We end this section by considering the specific case when $\beta = 0$, *i.e.*, when the channel is synchronous. For a given sampling frequency ρ , the receiver gets to see only a fraction ρ of the transmitted codeword (whether sampling is adaptive or non-adaptive) and hence

$$C(0, \rho) = \rho C(0, 1)$$

for any $\rho \geq 0$.

How is it possible that sparse output sampling induces a rate per unit cost loss for synchronous communication ($\beta = 0$), but not for asynchronous communication ($\beta > 0$) as we saw in Theorems 2 and 3? The reason for this is that when $\beta > 0$, the level of asynchronism is exponential in B . Therefore, even if the receiver is constrained to sample only a fraction ρ of the channel outputs, it may still occasionally sample fully over, say, $\Theta(B)$ channel outputs, and still satisfy the overall constraint that the fraction shouldn't exceed ρ .⁴

Remark 1. Theorems 2 and 3 remain valid under universal decoding, *i.e.*, the only element from the channel that the decoder needs to know is its output alphabet \mathcal{Y} . This is briefly discussed at the end of Section IV.

Remark 2. Consider a multiple access generalization of the point-to-point setting where, instead of one transmitter, there are

$$U = 2^{vB}$$

transmitters who communicate to a common receiver, where v , $0 \leq v \leq \beta$, denotes the occupation parameter of the channel. The messages arrival times $\{\nu_1, \nu_2, \dots, \nu_U\}$ at the transmitters are jointly independent and uniformly distributed over $[1, \dots, A]$ with $A = 2^{\beta B}$ as before. Communication takes place as in the previous point-to-point case, each user uses the same codebook, and transmissions start at the times $\{\sigma_1, \sigma_2, \dots, \sigma_U\}$. Whenever a user tries to access the channel while it is occupied, the channel outputs random symbols, independent of the input (collision model).

The receiver operates sequentially and declares U messages at the times

$$S_{\tau_1}, S_{\tau_1} + S_{\tau_2}, \dots, S_{\tau_1} + S_{\tau_2} + \dots S_{\tau_U}$$

where stopping time τ_i , $1 \leq i \leq U$, is with respect to the output samples

$$Y_{S_{\tau_{i-1}}}, Y_{S_{\tau_{i-1}}+1}, Y_{S_{\tau_{i-1}}+2}, \dots$$

It is easy to check (say, from the Birthday problem [12]) that if

$$v < \beta/2$$

and hence $U = o(\sqrt{A}) = o(2^{\beta B/2})$, the collision probability goes to zero as $B \rightarrow \infty$. Hence in the regime of large message size, the transmitters are (essentially) operating orthogonally, and each user can achieve the point-to-point capacity per unit cost assuming a per/user error probability. We may refer to this regime as the

⁴If over a long trip we have a high-mileage drive, we can still push the car a few times without impacting the overall mileage.

regime of “sparse transmissions,” relevant in a sensor network monitoring independent rare events.

Note that since the users use the same codebook, the receiver does not know which transmitter conveys what information. The receiver can only recognize the set of transmitted messages.

If the receiver is also required to identify the messages and their transmitters, then each transmitter effectively conveys $B(1 + \nu)$ information bits and the capacity per unit cost gets multiplied by $1/(1 + \nu)$.

IV. ANALYSIS

The following two standard type results are often used in our analysis.

Fact 1 ([10, Lemma 1.2.2]).

$$|\mathcal{P}_n^{\mathcal{X}}| = \text{poly}(n).$$

Fact 2 ([10, Lemma 1.2.6]). If X^n is independent and identically distributed (i.i.d.) according to $P_1 \in \mathcal{P}^{\mathcal{X}}$, then

$$\text{poly}(n)e^{-nD(P_2\|P_1)} \leq \mathbb{P}(X^n \in \mathcal{T}_{P_2}) \leq e^{-nD(P_2\|P_1)}.$$

for any $P_2 \in \mathcal{P}_n^{\mathcal{X}}$.

Achievability of Theorem 2: Fix some arbitrary distribution P on \mathcal{X} . Let X be the input having that distribution and let Y be the corresponding output, i.e., $(X, Y) \sim P(\cdot)Q(\cdot|\cdot)$.

Given B bits of information to be transmitted, the codebook \mathcal{C} is randomly generated as follows. For each message $m = 1, \dots, M$, randomly generate length n sequences x^n i.i.d. according to P , until x^n belongs to the “constant composition” set⁵

$$\mathcal{A}_n = \{x^n : \|\hat{P}_{x^n} - P\| \leq 1/\log n\}. \quad (3)$$

If (3) is satisfied, then let $c^n(m) = x^n$ and move to the next message. Stop when a codeword has been assigned to all messages. From Chebyshev’s inequality, for any fixed m , no repetition will be required with high probability to generate $c^n(m)$, i.e.,

$$P^n(\mathcal{A}_n) \rightarrow 1 \quad \text{as } n \rightarrow \infty \quad (4)$$

where P^n denotes the order n product distribution of P .

The obtained codewords are thus essentially of constant composition—i.e., each symbol appears roughly the same number of times—and have cost $n\mathbb{E}[k(X)](1 + o(1))$ as $n \rightarrow \infty$ where $k(\cdot)$ is the input cost function of the channel.

Case $\star \in \mathcal{X}$: Information transmission is as follows. For simplicity let us first assume that $1/\rho$ is an integer.

Codeword symbols can be transmitted only at multiples of $1/\rho$. Times that are integer multiples of $1/\rho$ from now on are referred to as transmission times. Given a message m available at time ν , the transmitter sends the corresponding codeword $c^n(m)$ during the first n information transmission times coming at time $\geq \nu$. In between transmission times the transmitter sends \star . Hence, the transmitter sends

$$c_1(m) \star \dots \star c_2(m) \star \dots \star c_3(m) \{ \dots \} c_n(m)$$

starting at time $\sigma = \sigma(\nu) = \min\{t \geq \nu : \lfloor t/\rho \rfloor \geq \nu\}$.

The receiver operates as follows. Sampling is performed only at the transmission times. At transmission time t , the decoder computes the empirical distributions

$$\hat{P}_{c^n(m), y^n}(\cdot, \cdot)$$

induced by the last output samples y^n and all the codewords $\{c^n(m)\}$. If there is a unique message m for which

$$\|\hat{P}_{c^n(m), y^n}(\cdot, \cdot) - P(\cdot)Q(\cdot|\cdot)\| \leq 2/\log n,$$

the decoder stops and declares that message m was sent. If two (or more) codewords $c^n(m)$ and $c^n(m')$ relative to two different messages m and m' are typical with y^n , the decoder stops and declares one of the corresponding messages at random. If no codeword is typical with y^n , the decoder repeats the procedure at the next transmission time. If by the time of the last transmission time no message has been declared, the decoder outputs a random message.

We first compute the error probability averaged over codebooks and messages. Suppose message m is transmitted. The error event that the decoder declares some specific message $m' \neq m$ can be decomposed as⁶

$$\{m \rightarrow m'\} = \mathcal{E}_1 \cup \mathcal{E}_2, \quad (5)$$

where the error events \mathcal{E}_1 and \mathcal{E}_2 are defined as

- \mathcal{E}_1 : the decoder stops at a time t between σ and $\sigma + (2n - 2)/\rho$ (including σ and $\sigma + (2n - 2)/\rho$) and declares m' ;
- \mathcal{E}_2 : the decoder stops either at a time t before time σ or from time $\sigma + (2n - 1)/\rho$ onwards and declares m' .

Note that when event \mathcal{E}_1 happens, the observed sequence is generated by the sent codeword. By contrast, when event \mathcal{E}_2 happens, then the observed sequence is generated only by pure noise.

Using analogous arguments as in the achievability of [1, Proof of Theorem 1] we obtain the upper bounds

$$\mathbb{P}_m(\mathcal{E}_1) \leq 2^{-n(I(X;Y) - \epsilon)}$$

⁵ $\|\cdot\|$ refers to the L_1 -norm.

⁶Notice that the decoder outputs a message with probability one.

and

$$\mathbb{P}_m(\mathcal{E}_2) \leq A \cdot 2^{-n(I(X;Y)+D(Y||Y_*)-\varepsilon)}$$

which are both valid for any fixed $\varepsilon > 0$ provided that n is large enough.

Combining, we get

$$\begin{aligned} \mathbb{P}_m(m \rightarrow m') &\leq 2^{-n(I(X;Y)-\varepsilon)} \\ &\quad + A \cdot 2^{-n(I(X;Y)+D(Y||Y_*)-\varepsilon)}. \end{aligned}$$

Hence, taking a union bound over all possible wrong messages, we obtain that for all $\varepsilon > 0$,

$$\begin{aligned} \mathbb{P}(\mathcal{E}) &\leq 2^B \left(2^{-n(I(X;Y)-\varepsilon)} \right. \\ &\quad \left. + A \cdot 2^{-n(I(X;Y)+D(Y||Y_*)-\varepsilon)} \right) \\ &\stackrel{\text{def}}{=} \varepsilon_1(n) \end{aligned} \quad (6)$$

for n large enough.

We now show that the delay of our coding scheme in the sense of Definition 4 is at most n/ρ . Suppose a specific (non-random) codeword $c^n(m) \in \mathcal{A}$ is sent. If

$$\tau > \sigma + (n-1)/\rho,$$

then necessarily $c^n(m)$ is not typical with $Y_\sigma^{\sigma+(n-1)/\rho}$. By Sanov's theorem this happens with vanishing error probability and hence

$$\mathbb{P}(\tau - \sigma \leq (n-1)/\rho) = 1 - \varepsilon_2(n)$$

with $\varepsilon_2(n) \rightarrow 0$ as $n \rightarrow \infty$. Hence, since $\nu \leq \sigma < \nu + 1/\rho$, we get

$$\mathbb{P}(\tau - \nu \leq n/\rho) = 1 - \varepsilon_2(n).$$

The proof can now be concluded. From inequality (6) there exists a specific code $\mathcal{C} \subset \mathcal{A}_n$ whose error probability, averaged over messages, is less than $\varepsilon_1(n)$. Removing the half of the codewords with the highest error probability, we end up with a set \mathcal{C}' of 2^{B-1} codewords whose maximum error probability $\mathbb{P}(\mathcal{E})$ is such that

$$\mathbb{P}(\mathcal{E}) \leq 2\varepsilon_1(n), \quad (7)$$

and whose delay satisfies

$$d(\mathcal{C}', \varepsilon_2(n)) \leq n/\rho.$$

Now fix the ratio B/n and substitute $A = 2^{\beta B}$ in the definition of $\varepsilon_1(n)$ (see (6)). Then, $\mathbb{P}(\mathcal{E})$ goes to zero as $B \rightarrow \infty$ whenever

$$\frac{B}{n} < \min \left\{ I(X;Y), \frac{I(X;Y) + D(Y||Y_*)}{1 + \beta} \right\}. \quad (8)$$

Recall that by construction, all the codewords have cost $n\mathbb{E}[k(X)](1 + o(1))$ as $n \rightarrow \infty$. Hence, for any $\eta > 0$ and all n large enough

$$k(\mathcal{C}') \leq n\mathbb{E}[k(X)](1 + \eta). \quad (9)$$

Condition (8) is thus implied by condition

$$\frac{B}{K(\mathcal{C}')} < \min \left\{ \frac{I(X;Y)}{(1 + \eta)\mathbb{E}[k(X)]}, \frac{I(X;Y) + D(Y||Y_*)}{\mathbb{E}[k(X)](1 + \eta)(1 + \beta)} \right\}. \quad (10)$$

Maximizing over all input distributions and using the fact that $\eta > 0$ is arbitrary proves that $\mathcal{C}(\beta, 1)$ —where $\mathcal{C}(\beta, 1)$ is defined in Theorem 1—is asymptotically achieved by non-random codes with delay no larger than n/ρ with probability approaching one as $n \rightarrow \infty$.

Finally, if $1/\rho$ is not an integer, it suffices to define transmission times as

$$t_j = \lfloor j/\rho \rfloor.$$

This guarantees the same asymptotic performance as for the case where $1/\rho$ is an integer.

Case $\star \notin \mathcal{X}$: Parse the entire sequence $\{1, 2, \dots, A + n - 1\}$ into consecutive superperiods of size n/ρ —take $\lfloor n/\rho \rfloor$ if n/ρ is not an integer. The periods of duration n occurring at the end of each superperiod are referred to as transmission periods. Given ν , the codeword starts being sent over the first transmission period starting at a time $> \nu$. In particular, if ν happens over a transmission period, then the transmitter delays the codeword transmission to the next superperiod.

The receiver sequentially samples only the transmission periods. At the end of a transmission period, the decoder computes the empirical distributions

$$\hat{P}_{c^n(m), y^n}(\cdot, \cdot)$$

induced by the last output samples y^n and all the codewords $\{c^n(m)\}$. If there is a unique message m for which

$$\|\hat{P}_{c^n(m), y^n}(\cdot, \cdot) - P(\cdot)Q(\cdot)\| \leq 2/\log n,$$

the decoder stops and declares that message m was sent. If two (or more) codewords $c^n(m)$ and $c^n(m')$ relative to two different messages m and m' are typical with y^n , the decoder stops and declares one of the corresponding messages at random. If no codeword is typical with y^n , the decoder waits for the next transmission period to occur, samples it, and repeats the decoding procedure. Similarly as for the previous case, if at the end of the last transmission period no message has been declared, the decoder outputs a random message.

Following the same arguments as for the case $\star \in \mathcal{X}$ we deduce that (10) also holds in this case and that for the delay we have

$$\mathbb{P}(\tau - \nu \leq n + n/\rho) = 1 - \varepsilon_2(n)$$

for some $\varepsilon_2(n) \rightarrow 0$ as $n \rightarrow \infty$. To see this, note that a superperiod has duration n/ρ and that if ν happens during a transmission period, then the actual codeword transmission is delayed to the next transmission period. ■

Delay Converse of Theorem 2: We consider the cases $\star \in \mathcal{X}$ and $\star \notin \mathcal{X}$ separately.

Case $\star \in \mathcal{X}$: Pick some arbitrary $0 < \rho < 1$, $\beta > 0$ such that $\mathcal{C}(\beta, \rho) > 0$, and $0 < \gamma < 1$. Consider a code \mathcal{C}_B with length n_B codewords that achieves rate per unit cost $\gamma \mathcal{C}(\beta, \rho) - \varepsilon_B > 0$, maximum error probability at most ε_B , sampling frequency $\rho(\mathcal{C}_B, \varepsilon_B) \leq \rho + \varepsilon_B$, and delay $d_B = d(\mathcal{C}_B, \varepsilon_B)$, for some $\varepsilon_B \xrightarrow{B \rightarrow \infty} 0$. The sampling strategy \mathcal{S} is supposed to be non-adaptive, and for the moment also non-randomized.

Denote by \mathcal{J}_γ the event that the decoder samples at least γn_B^* samples of the sent codeword—recall that n_B^* refers to the minimal codeword length, see Section III. Then by the converse of the [1, Theorem 1]

$$\mathbb{P}_m(\mathcal{J}_{\gamma'}) = 1 - o(1) \quad (B \rightarrow \infty) \quad (11)$$

for any message m , where $\gamma' = \gamma'(\gamma)$ satisfies $\gamma' = \gamma'(\gamma) > 0$ for any $\gamma > 0$ and $\lim_{\gamma \rightarrow 1} \gamma'(\gamma) = 1$.

Further, by our assumption on the error probability and on the delay (see Definition 4), we have for any message m

$$\mathbb{P}_m(\mathcal{E}_m^c \cap \{\tau - \nu \leq d_B - 1\}) \rightarrow 1 \quad (B \rightarrow \infty),$$

where \mathcal{E}_m^c denotes the successful decoding event. This implies that for any message m

$$\mathbb{P}_m(0 \leq \tau - \nu \leq d_B - 1) \rightarrow 1 \quad (B \rightarrow \infty),$$

since the error probability is bounded away from zero whenever $\tau < \nu$.

It then follows that

$$\mathbb{P}_m(\{0 \leq \tau - \nu \leq d_B - 1\} \cap \mathcal{J}_{\gamma'}) = 1 - o(1) \quad (B \rightarrow \infty). \quad (12)$$

Hence, since ν is uniformly distributed over $\{1, 2, \dots, A + n - 1\}$, for B large enough we have

$$\mathbb{P}_m(\{0 \leq \tau - t \leq d_B - 1\} \cap \mathcal{J}_{\gamma'} | \nu = t) > 0$$

for at least $(1 - o(1))A$ values of $t \in \{1, 2, \dots, A\}$. Now, conditioned on $\{\nu = t\}$, if event

$$\{0 \leq \tau - t \leq d_B - 1\} \cap \mathcal{J}_{\gamma'}$$

happens (*i.e.*, with non-zero probability), then necessarily the period $\{t, t + 1, \dots, t + d_B - 1\}$ contains at least $\gamma' n_B^*$ sampling times—here we use the fact that \mathcal{S} is non-randomized.

It then follows that

$$|\mathcal{S}| \geq \left\lfloor \frac{(1 - o(1))A}{d_B} \right\rfloor \gamma' \cdot n_B^*. \quad (13)$$

Now if

$$\rho d_B \leq n_B^*(1 - \varepsilon) \quad (14)$$

for some arbitrary fixed $0 < \varepsilon < 1$, then

$$\left\lfloor \frac{(1 - o(1))A}{d_B} \right\rfloor \gamma' n_B^* \geq \frac{(1 - o(1))\gamma'}{1 - \varepsilon} \rho A (1 - o(1)) \quad (15)$$

as $B \rightarrow \infty$.

Hence, by taking γ' and hence γ close enough to 1 and by taking B large enough

$$(1 - o(1))\gamma'/(1 - \varepsilon) > 1.$$

Therefore, if (14) holds, from (13) and (15) we get

$$|\mathcal{S}| \geq \rho(1 + \varepsilon')A \quad (16)$$

for B large enough and some $\varepsilon' > 0$ such that $\varepsilon' \rightarrow 0$ as $\varepsilon \rightarrow 0$. Inequality (16) implies that the sampling constraint is violated, as we now show.

Fix an arbitrary $0 < \varepsilon'' < 1$. For an arbitrary integer $1 \leq k \leq A + n - 1$ and any message m

$$\begin{aligned} \mathbb{P}_m(S_\tau \geq \rho\tau(1 + \varepsilon'')) &\geq \mathbb{P}_m(S_\tau \geq \rho\tau(1 + \varepsilon'') | \tau \geq k) \mathbb{P}_m(\tau \geq k) \\ &\geq \mathbb{P}_m(S_k \geq \rho(A + n - 1)(1 + \varepsilon'') | \tau \geq k) \mathbb{P}_m(\tau \geq k) \\ &= \mathbb{P}_m(S_k \geq (1 + \varepsilon'')\rho A(1 + o(1))) \mathbb{P}_m(\tau \geq k) \end{aligned} \quad (17)$$

where for the second inequality we used the fact that the sampling times S_1, S_2, \dots are non-decreasing and the fact that $\tau \leq A + n - 1$. We now show that both terms

$$\mathbb{P}_m(S_k \geq (1 + \varepsilon'')\rho A(1 + o(1)))$$

and

$$\mathbb{P}_m(\tau \geq k)$$

are bounded away from zero in the limit $B \rightarrow \infty$, for an appropriate choice of k . This, by (17), implies that

$$\liminf_{B \rightarrow \infty} \mathbb{P}_m(S_\tau \geq \rho\tau(1 + \varepsilon'')) > 0,$$

i.e., that sampling frequency ρ is not achievable whenever (14) holds. In other words, to achieve a sampling

frequency ρ it is necessary that delay and codeword length satisfy

$$d_B \geq \frac{n_B^*}{\rho}(1 - o(1)).$$

Let

$$k = (1 + 2\varepsilon'')\rho A. \quad (18)$$

Since $S_k \geq k$,

$$S_k \geq (1 + 2\varepsilon'')\rho A,$$

and so by choosing $\varepsilon'' > 0$ small enough we get

$$\mathbb{P}(S_k \geq (1 + \varepsilon'')\rho A(1 + o(1))) = 1 \quad (19)$$

for B large enough.

Since \mathcal{C}_B achieves (maximum) error probability $\leq \varepsilon_B$ we have for any message m

$$\begin{aligned} \varepsilon_B &\geq \mathbb{P}_m(\mathcal{E}) \\ &\geq \mathbb{P}_m(\mathcal{E} | \tau < k, \nu \geq k) \mathbb{P}_m(\tau < k, \nu \geq k) \\ &\geq \frac{1}{2} \mathbb{P}_m(\tau < k | \nu \geq k) \mathbb{P}(\nu \geq k) \\ &= \frac{1}{2} \mathbb{P}_*(\tau < k) \mathbb{P}(\nu \geq k). \end{aligned} \quad (20)$$

For the third inequality in (20) note that event $\{\tau < k, \nu \geq k\}$ means that the decoder declares a message before the actual message even starts being sent. In this case, the error probability is at least $1/2$, since a message set always contains at least two messages (see Section II). For the last equality in (20), note that event $\{\tau \geq k\}$ depends only on Y_1^k , which are i.i.d. $\sim Q_*$ when conditioned on $\{\nu > k\}$ — \mathbb{P}_* denotes the output distribution under pure noise, i.e., when Y_1^{A+n-1} is an i.i.d. Q_* random sequence. Repeating this last change of measure argument we get

$$\begin{aligned} \mathbb{P}_m(\tau \geq k) &\geq \mathbb{P}_m(\tau \geq k | \nu > k) \mathbb{P}(\nu > k) \\ &= \mathbb{P}_*(\tau \geq k) \mathbb{P}(\nu > k) \\ &\geq (1 - 2\varepsilon_B / \mathbb{P}(\nu \geq k)) \mathbb{P}(\nu > k) \\ &= (1 - o(1)) \mathbb{P}(\nu > k) \\ &= (1 - \rho(1 + 2\varepsilon''))(1 - o(1)) \quad B \rightarrow \infty. \end{aligned} \quad (21)$$

The second inequality follows from (20). For the second and third equality in (21) we use the fact that ν is uniformly distributed over $\{1, 2, \dots, A\}$, hence by (18)

$$\mathbb{P}(\nu > k) = (A - k)/A = 1 - (1 - \rho(1 + 2\varepsilon'')) > 0.$$

Since $\rho(1 + \varepsilon'') > 0$, we have $\liminf_{B \rightarrow \infty} \mathbb{P}_m(\tau \geq k) > 0$ by (21), yielding the desired claim.

Finally, to see that randomized sampling strategies cannot achieve a better sampling frequency, note that a

randomized sampling strategy can be viewed as a probability distribution over deterministic sampling strategies. Therefore, because the previous analysis holds for any deterministic sampling strategy, it must also hold for randomized sampling strategies rules.

Case $\star \notin \mathcal{X}$: Pick some arbitrary $\varepsilon > 0$ and consider a code \mathcal{C}_B with length n_B codewords that achieves rate per unit cost $\mathcal{C}(\beta, \rho) - \varepsilon_B$, error probability $\leq \varepsilon_B$, delay $d_B = d(\mathcal{C}_B, \varepsilon_B)$, and sampling frequency $\rho(\mathcal{C}_B, \varepsilon_B) \leq \rho + \varepsilon_B$ for some $\varepsilon_B \xrightarrow{B \rightarrow \infty} 0$. As in the previous case, without loss of optimality the sampling strategy \mathcal{S} is supposed to be non-randomized.

Because $\star \notin \mathcal{X}$, we have $k(x) > 0$ for any $x \in \mathcal{X}$ and therefore to achieve the full-sampling asynchronous capacity per unit cost it is necessary that the codeword length remains essentially the same as under full sampling. More specifically, we must have

$$n_B \leq n'_B(1 + \eta(\varepsilon)) \quad B \rightarrow \infty \quad (22)$$

for some $\eta(\varepsilon) \rightarrow 0$ as $\varepsilon \rightarrow 0$, where n'_B denotes the number of sampled codeword positions—recall that codeword positions are the positions from time σ up to time $\sigma + n_B - 1$. Note that this is in contrast with the case $\star \in \mathcal{X}$, where the codeword transmission duration can be expanded by transmitting \star at no cost.

Proceeding as for the case $\star \in \mathcal{X}$, we have

$$\mathbb{P}_m(\{0 \leq \tau - t \leq d_B - 1\} \cap \mathcal{A}_{\gamma(\varepsilon)}) = 1 - o(1) \quad (23)$$

$$\mathbb{P}_m(\{0 \leq \tau - t \leq d_B - 1\} \cap \mathcal{A}_{\gamma(\varepsilon)} | \nu = t) > 0$$

for any $t \in \mathcal{B}$ where \mathcal{B} is a certain subset of $\{1, 2, \dots, A\}$ with $|\mathcal{B}| = (1 - o(1))A$. This means that for any $t \in \mathcal{B}$ the decoder samples a “block” $b \subseteq \mathcal{S}$ of cardinality at least γn_B over the period $[t, t+1, \dots, t+d_B-1]$. Moreover, if we denote by $i(b)$ and $f(b)$ the time position within $\{1, 2, \dots, A+n-1\}$ of the first and the last element of b , respectively, then for each block we have $f(b) - i(b) \leq n_B$.

Because of the sampling constraint, there are at most

$$N = \frac{\rho A(1 + o(1))}{\gamma n_B}$$

distinct blocks of size γn_B . This implies that d_B should satisfy

$$d_B \geq (\gamma n_B / \rho + \gamma n_B)(1 - o(1)),$$

as we now show. Intuitively, the reason the delay must satisfy this bound is that because the codewords must now be blocks of symbols, the receiver might as well sample in blocks of n_B symbols. Then, the sampling

constraint means that, on average, the gap between sampled blocks grows like n_B/ρ . However, if the message arrives at a time ν close to the beginning of a block, then in addition to waiting until the next block, the message must wait for most of the current block before being transmitted—close to capacity we cannot afford to miss a portion of the codeword other than negligible. Therefore, the delay must grow as $n_B/\rho + n_B$. We formalize this reasoning below.

Suppose, by way of contradiction, that

$$d_B \leq \gamma n_B(1 + 1/\rho)(1 - \varepsilon) \quad (24)$$

for some $\varepsilon > 0$, and assume for the moment that each block $b_{(t)}$ is composed of $\gamma n_B(1 + o(1))$ elements and that there are least $N(1 - o(1))$ distinct blocks.

Define the “occupation” slot of a block as γn_B plus the time interval until the next block. The average occupation slot per block is thus

$$\frac{A}{N} = \frac{\gamma n_B}{\rho}(1 - o(1)).$$

Hence, for any $\varepsilon' > 0$ there is a set of at least ηN occupation slots each of size at most

$$\frac{\gamma n_B}{\rho}(1 + \varepsilon')$$

where $\eta = \eta(\varepsilon') > 0$ for any $\varepsilon' > 0$. Consider such a set of occupation slots for some $\varepsilon' > 0$ which is specified later, let b be a block belonging to one such slots, and let b' denote the block coming after b .⁷ Denote by $i(b)$ and $f(b)$ the time position within $\{1, 2, \dots, A + n - 1\}$ of the first and the last element of b , respectively. Then for B large enough

$$f(b') - i(b) \geq \gamma n_B(1 + \frac{1}{\rho}(1 + \varepsilon'))$$

and therefore by taking $\varepsilon' > 0$ small enough we get

$$f(b') - i(b) > \eta' n_B$$

by (24), where $\eta' = \eta'(\varepsilon, \varepsilon') > 0$ for any $\varepsilon > 0$ and $\varepsilon' > 0$. It then follows that for any $t \in (i(b), i(b) + \eta' n_B]$, the interval $[t, t + 1, \dots, t + d_B - 1]$ contains neither blocks b and b' completely. Therefore, conditioned on $\nu \in (i(b), i(b) + \eta' n_B]$, event

$$\{\tau - \nu \leq d_B - 1\} \cap \mathcal{J}_\gamma$$

does not happen. It then follows that if (24) holds for some $\varepsilon > 0$, then

$$\limsup_{B \rightarrow \infty} \mathbb{P}_m(\{\tau - \nu \leq d_B - 1\} \cap \mathcal{J}_\gamma) \leq 1 - \eta\eta' < 1$$

⁷For reasons that will soon be obvious, b should not be the right most block within the set.

which contradicts (23).

The above argument assumes that there are N disjoint blocks of size γn_B . If there are fewer and possibly larger blocks, the arguments easily extend by defining the blocks b as any subset of \mathcal{S} such that $f(b) - i(b) \leq n_B$. ■

Proof of Theorem 3: We show that $C(\beta, \rho) = C(\beta, 1)$ for any $\beta > 0$ and $0 < \rho \leq 1$ and that $C(\beta, 1)$ can be achieved with codes $\{\mathcal{C}_B\}$ with delay $d(\mathcal{C}_B, \varepsilon_B) = n_B^*(1 + o(1))$ as $B \rightarrow \infty$.

Let P be the distribution achieving $C(\beta, 1)$ (see Theorem 1). We generate 2^B codewords of length

$$n - \log(n)$$

as in the proof of Theorem 2 according to distribution P . Each codeword starts with a common preamble that consists of $\log(n)$ repetitions of a symbol x such that $Q(\cdot|x) \neq Q(\cdot|\star)$.

For the proposed asymptotically optimal sampling/decoding strategy, it is convenient to introduce the following notation. Let \tilde{Y}_a^b denote the random vector obtained by extracting the components of the output process Y_t at $t \in [a, b]$ of the form $t = \lceil j/\rho \rceil$ for non-negative integer j . Notice that, for any $t \geq \ell$ and $\ell \gg 1$, $\tilde{Y}_{t-\ell+1}^t$ contains $\approx \rho\ell$ samples.

The strategy starts in the sparse mode, taking samples at times $S_j = \lceil j/\rho \rceil$, $j = 1, 2, \dots$. At each j , the receiver computes the empirical distribution (or type)

$$\hat{P}_j = \hat{P}_{\tilde{Y}_{S_j - \log(n) + 1}^{S_j}}$$

of the sampled output in the most recent window of length $\log(n)$.

If the probability of this type under pure noise is large enough, i.e., if

$$\mathbb{P}_*(\mathcal{T}_{\hat{P}_j}) > \frac{1}{n^2},$$

the mode is kept unchanged and we repeat this test at the next round $j + 1$.

Instead, if

$$\mathbb{P}_*(\mathcal{T}_{\hat{P}_j}) \leq \frac{1}{n^2},$$

then we switch to the dense sampling mode, taking samples continuously for at most n time steps. At each of these steps the receiver applies the same sequential typicality decoder as in the proof of Theorem 2, based on the past $n - \log n$ output samples. If no codeword is typical with the channel outputs after these n times steps, sampling is switched back to the sparse mode.

We compute the error probability of the above scheme, its relative number of samples, and its delay.

For the error probability, a similar analysis as for the non-adaptive case in the proof of Theorem 2 still applies,

with ρn being replaced by $n - \log n$. In particular, after fixing the ratio B/n and thereby imposing a delay linear in B , equation (10) holds with $\rho = 1$.

For the relative number of samples, we now show that

$$\mathbb{P}_m(\tau/S_\tau \geq \rho + \varepsilon_B) \xrightarrow{n \rightarrow \infty} 0 \quad (25)$$

with $\varepsilon_B = 1/\text{poly}(B)$ from which we then conclude that $\mathcal{C}(\beta, \rho) \geq \mathcal{C}(\beta, 1)$. To do this, it is convenient to introduce Z_i , $1 \leq i \leq A + n - 1$, which is equal to one if at time i the receiver switches to the dense mode and samples the next n channel outputs and equal to zero otherwise. Then it follows that

$$\tau \leq \rho S_\tau + n \sum_{i=1}^{S_\tau} Z_i. \quad (26)$$

To see this, note that the number of samples involved in the sparse mode is at most ρS_τ and that the number of samples involved in the dense mode is at most $n \sum_{i=1}^{S_\tau} Z_i$ (it is actually equal to $n \sum_{i=1}^{S_\tau} Z_i$ if we ignore the boundary discrepancies that we cannot sample beyond time $A + n - 1$).

From (26)

$$\begin{aligned} \mathbb{P}(\tau/S_\tau \geq \rho + \varepsilon) &\leq \mathbb{P}(n \sum_{i=1}^{S_\tau} Z_i \geq S_\tau \varepsilon) \\ &\leq \mathbb{P}(n \sum_{i=1}^{S_\tau} Z_i \geq S_\tau \varepsilon, \nu \leq S_\tau \leq \nu + 2n - 2) \\ &\quad + \mathbb{P}(S_\tau < \nu \text{ or } S_\tau > \nu + 2n - 2). \end{aligned} \quad (27)$$

We now show that the right-hand side of the second inequality in (27) vanishes as $B \rightarrow \infty$.

For the first term on the right-hand side of the second inequality in (27), since the Z_i 's are nonnegative

$$\begin{aligned} \mathbb{P}(n \sum_{i=1}^{S_\tau} Z_i \geq S_\tau \varepsilon; \nu \leq S_\tau \leq \nu + 2n - 2) \\ \leq \mathbb{P}(n \sum_{i=1}^{\nu+2n-2} Z_i \geq \nu \varepsilon). \end{aligned} \quad (28)$$

Now, conditioned on $\nu = t$, the Z_i 's, $1 \leq i \leq t - 1$, are binary random variables distributed according to pure noise. Hence,

$$\begin{aligned} \mathbb{P}\left(n \sum_{i=1}^{t+2n-2} Z_i \geq t\varepsilon \mid \nu = t\right) \\ \leq \mathbb{P}_\star\left(n \sum_{i=1}^{t-1} Z_i \geq t\varepsilon - (2n - 1)\right) \\ \leq \frac{t - 1}{(t\varepsilon - (2n - 1) - (t - 1)/n^2)^2} \\ = o(1) \quad (t \rightarrow \infty) \end{aligned} \quad (29)$$

where the second inequality follows from Chebyshev's inequality and by noting that for $1 \leq i \leq t - 1$ we have

$$\text{Var}(Z_i) \leq \mathbb{E}Z_i \leq 1/n^2$$

since the variance of a Bernoulli random variable is at most its mean which, in turn, is at most $1/n^2$.

Therefore,

$$\begin{aligned} \mathbb{P}\left(n \sum_{i=1}^{\nu+2n-1} Z_i \geq \nu \varepsilon\right) \\ \leq \mathbb{P}(\nu \leq \sqrt{A}) \\ + \frac{1}{A} \sum_{t=\sqrt{A}+1}^A \mathbb{P}(n \sum_{i=1}^{\nu+2n-1} Z_i \geq \nu \varepsilon \mid \nu = t) \\ = o(1) \quad (B \rightarrow \infty) \end{aligned} \quad (30)$$

where the last equality follows from (29) and the fact that ν is uniformly distributed over $\{1, 2, \dots, A = e^{\beta B}\}$. From (28) and (30) we get

$$\mathbb{P}(n \sum_{i=1}^{S_\tau} Z_i \geq S_\tau \varepsilon; \nu \leq S_\tau \leq \nu + 2n - 2) = o(1)$$

as $B \rightarrow \infty$.

We now show that

$$\mathbb{P}(S_\tau < \nu \text{ or } S_\tau \geq \nu + 2n - 1) \rightarrow 0 \quad (B \rightarrow \infty). \quad (31)$$

That $\mathbb{P}(S_\tau < \nu) \rightarrow 0$ follows from the fact that $\mathbb{P}_m(\mathcal{E}_2) \rightarrow 0$ where \mathcal{E}_2 is defined in the proof of Theorem 2. That $\mathbb{P}(S_\tau \geq \nu + 2n - 1) \rightarrow 0$ follows from the fact that with probability tending to one the sampling strategy will change mode over the transmitted codeword and that the typicality decoder will make a decision up to time $\nu + n - 1$ with probability tending to one. This last argument can also be used for the delay to show that $d(\mathcal{C}_B, \varepsilon_B) = n(1 + o(1))$ for some $\varepsilon_B \rightarrow 0$.

Finally, by optimizing the input distribution to minimize delay (see paragraph after Theorem 1) we deduce that $\mathcal{C}(\beta, 1) = \mathcal{C}(\beta, \rho)$ and that the capacity per unit cost is achievable with delay $n_B^*(1 + o(1))$. ■

We end this section with a few words concerning the Remark 1 at the end of Section III. To prove the claim it suffices to slightly modify the achievability schemes yielding Theorems 2 and 3 to make them universal at the decoder.

The first modification is needed to estimate the pure noise distribution Q_\star with a negligible fraction of channel outputs. An estimate of this distribution is obtained by sampling the first \sqrt{A} output symbols. At the end of this estimation phase, the receiver declares the pure noise

distribution as being equal to $\hat{P}_{Y_1^{\sqrt{A}}}$. Note that since ν is uniformly distributed over $\{1, 2, \dots, A\}$ we have

$$\mathbb{P}(\|\hat{P}_{Y_1^{\sqrt{A}}} - Q_\star\|_1 \geq \varepsilon_B) \rightarrow 1$$

as $B \rightarrow \infty$, for some $\varepsilon_B \rightarrow 0$. Note also that this estimation phase requires a negligible amount of sampling, *i.e.*, sublinear in A .

The second modification concerns the typicality decoder which is replaced by an MMI (Maximum Mutual Information) decoder (see [10, Chapter 2]).

It is straightforward to verify that the modified schemes indeed achieve the asynchronous capacity per unit cost. The formal arguments are similar to those used in [3, Proof of Theorem 2] (see also [5, Theorem 3] which proves the claim under full sampling and unit input cost) and are thus omitted.

REFERENCES

- [1] V. Chandar, A. Tchamkerten, and D. Tse, “Asynchronous capacity per unit cost,” *Information Theory, IEEE Transactions on*, vol. 59, no. 3, pp. 1213–1226, march 2013.
- [2] V. Chandar, A. Tchamkerten, and G. Wornell, “Optimal sequential frame synchronization,” *Information Theory, IEEE Transactions on*, vol. 54, no. 8, pp. 3725–3728, 2008.
- [3] A. Tchamkerten, V. Chandar, and G. Wornell, “Communication under strong asynchronism,” *Information Theory, IEEE Transactions on*, vol. 55, no. 10, pp. 4508–4528, 2009.
- [4] A. Tchamkerten, V. Chandar, and G. W. Wornell, “Asynchronous communication: Capacity bounds and suboptimality of training,” *Information Theory, IEEE Transactions on*, vol. 59, no. 3, pp. 1227–1255, march 2013.
- [5] Y. Polyanskiy, “Asynchronous communication: Exact synchronization, universality, and dispersion,” *Information Theory, IEEE Transactions on*, vol. 59, no. 3, pp. 1256–1270, march 2013.
- [6] D. Wang, V. Chandar, S. Chung, and G. Wornell, “Error exponents in asynchronous communication,” in *Information Theory Proceedings (ISIT), 2011 IEEE International Symposium on*. IEEE, 2011, pp. 1071–1075.
- [7] I. Shomorony, R. Etkin, F. Parvaresh, and A. Avestimehr, “Bounds on the minimum energy-per-bit for bursty traffic in diamond networks,” in *Information Theory Proceedings (ISIT), 2012 IEEE International Symposium on*. IEEE, 2012, pp. 801–805.
- [8] M. Khoshnevisan and J. Laneman, “Achievable rates for intermittent communication,” in *Information Theory Proceedings (ISIT), 2012 IEEE International Symposium on*, july 2012, pp. 1346–1350.
- [9] R. L. Dobrushin, “Asymptotic bounds on the probability of error for the transmission of messages over a memoryless channel using feedback,” *Probl. Kibern.*, vol. 8, pp. 161–168, 1963.
- [10] I. Csiszár and J. Körner, *Information Theory: Coding Theorems for Discrete Memoryless Channels*. New York: Cambridge University Press, 2011.
- [11] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms, 2nd edition*. MIT Press, McGraw-Hill Book Company, 2000.
- [12] G. Grimmett and D. Stirzaker, *One Thousand Exercises in Probability*. New York: Oxford University Press, 2001.